

1 **METHOD FOR PREDICTIVE POWER MANAGEMENT FOR OPERATING A DISK**
2 **DRIVE IN A MOBILE DEVICE TO OPTIMIZE POWER USAGE**

3 **BACKGROUND OF THE INVENTION**

4 **Field of the Invention**

5 The present invention relates to digital data storage, and more particularly, to operation of
6 a hard disk drive in a mobile device to optimize power usage.

7 **Description of the Prior Art**

8 A handheld mobile device has limited battery capacity and power consumption of each
9 component within the device may be of critical concern. A hard disk drive provides digital data
10 storage on a magnetic media of a rotating disk. A spindle motor rotates the disk at a
11 predetermined rate during access to the disk. The power consumption of the spindle motor may
12 be reduced by “spinning down” the disk between disk accesses thus allowing removal of a power
13 consuming control current to the spindle motor. However, spinning down the disk affects the
14 response time of the disk drive to disk commands and may involve a power consumption penalty
15 associated with “spinning up” the disk drive.

16 Accordingly, there exists a need for a technique for operating a disk drive in a mobile
17 device to optimize power usage. The present invention satisfies these needs.

18
19 **SUMMARY OF THE INVENTION**

20 The present invention may be embodied in a mobile device having a disk drive, and a
21 method for operating the disk drive in the mobile device to optimize power usage. Immediately
22 after responding to a disk access command from the mobile device, the disk drive is in a first
23 operating mode that consumes a relatively high level of power and may transition to a second
24 operating mode that consumes a relatively low level of power. In the method, a first time period
25 threshold and a second time period threshold may be set with the first time period threshold
26 being less than the second time period threshold. An adjustable delay time interval may be set to
27 correspond to one of the time period thresholds. The delay time interval is a time interval for
28 waiting after an end of a response to a disk access command before transitioning from the first
29 operating mode to the second operating mode. The disk drive measures a demand time interval

1 for each of a plurality of disk access commands. Each demand time interval may be defined as a
2 time period between an end of a response to a last disk access command and an arrival of a next
3 disk access command. If a majority of a predetermined number of the plurality of demand time
4 intervals falls within a time period less than the first time period threshold or a time period
5 greater than the second time period threshold, then the delay time interval may be set to be equal
6 to about the first time period threshold. If a majority of a predetermined number of the plurality
7 of demand time intervals falls within a time period between the first time period threshold and
8 the second time period threshold, the delay time interval may be set to be equal to about the
9 second time period threshold. The disk drive may transition from the first operating mode to the
10 second operating mode after expiration of the delay time interval.

11 In more detailed features of the invention, transitioning from the second operating mode
12 to the first operating mode may consume a transition quantity of power, and an exchange time
13 interval may be defined as a time period during which power consumed by operating in the first
14 power operating mode is about equal to the transition quantity of power. The predetermined
15 number of the plurality of demand time intervals may be 7 and the majority for setting the delay
16 time interval may be 4. Also, the delay time intervals comprising the predetermined number may
17 be measured using the most recent disk access commands.

18 In another embodiment of the invention, a third time period threshold may be set that is
19 greater than the second time period threshold and less than the exchange time interval. If a
20 majority of the predetermined number of the plurality of demand time intervals falls within a
21 time period less than the first time period threshold or a time period greater than the third time
22 period threshold, then the delay time interval may be set to be equal to about the first time period
23 threshold. If a majority of the predetermined number of the plurality of demand time intervals
24 falls within a time period between the first time period threshold and the second time period
25 threshold, the delay time interval may be set to be equal to about the second time period
26 threshold. If a majority of the predetermined number of the plurality of demand time intervals
27 falls within a time period between the second time period threshold and the third time period
28 threshold, then the delay time interval may be set to be equal to about the third time period
29 threshold. The first time period threshold may be about 0.3% of the exchange time interval, the
30 second time period threshold may be about 33% of the exchange time period, and the third time

1 period threshold may be about 66% of the exchange time interval.

2
3 **BRIEF DESCRIPTION OF THE DRAWINGS**

4 The accompanying drawings illustrate embodiments of the present invention and,
5 together with the description, serve to explain the principles of the invention.

6 FIG. 1 is a flow diagram of a method for operating a disk drive in a mobile device for
7 optimizing power usage using predictive power management, according to the present invention.

8 FIG. 2 is a timing diagram of demand time intervals with respect to two time period
9 thresholds and an adjustable delay time interval, according to the present invention.

10 FIG. 3 is a schematic diagram of a data structure having three bins for binning and
11 counting occurrences of demand time intervals, according to the present invention.

12 FIG. 4 is a block diagram showing a mobile device having a disk drive for implementing
13 the method of FIG. 1.

14 FIG. 5 is a graph diagram showing power consumption associated with power operating
15 modes of the disk drive of FIG. 4.

16 FIG. 6 is a timing diagram of demand time intervals with respect to three time period
17 thresholds and an adjustable delay time interval, according to the present invention.

18 FIG. 7 is a schematic diagram of a data structure having five bins for binning and
19 counting occurrences of demand time intervals, according to the present invention.

20
21 **DETAILED DESCRIPTION**

22 With reference to FIGS. 1-4, the present invention may be embodied in a method for
23 operating a disk drive 10 in a mobile device 12 to optimize power usage. The mobile device may
24 be a hand held computing device such as a personal digital assistant (PDA) or a mobile terminal
25 such as a cellular telephone. Immediately after responding to a disk access command, the disk
26 drive is in a first operating mode that consumes a relatively high level of power and may
27 transition to a second operating mode that consumes a relatively low level of power. A transition
28 from the second operating mode to the first operating mode consumes a transition quantity of
29 power. A predictive technique of the invention tracks a user's usage pattern and adapts an
30 operating mode transition time based on the usage pattern.

In the method, a first time period threshold TH1 and a second time period threshold TH2 are set (step 14). The first time period threshold is less than the second time period threshold. An adjustable delay time interval 16 is provided that may be set to correspond to one of the time period thresholds (step 18). The delay time interval is a time interval for waiting after an end of a response to a disk access command before transitioning from the first operating mode to the second operating mode. To set the adjustable delay time period, the disk drive services disk access commands and measures a demand time interval DMT(1-7) for each of a plurality of disk access commands (step 22). Each demand time interval is defined as a time period between an end of a response to a last disk access command and an arrival of a next disk access command of the predetermined number of disk access commands. If a majority of a predetermined number of the plurality of demand time intervals falls within a time period less than the first time period threshold or a time period greater than the second time period threshold (step 24), then the delay time interval is set to be equal to about the first time period threshold (step 26). However, if a majority of a predetermined number of the plurality of demand time intervals falls within a time period between the first time period threshold and the second time period threshold (step 24), the delay time interval is set to be equal to about the second time period threshold (step 28). The disk drive transitions from the first operating mode to the second operating mode after expiration of the newly set delay time interval (step 30) to optimize the disk drive's power usage. After another disk access command is received, a demand time interval is measured (step 22) and the delay time interval setting process repeats.

The time period less than the first time period threshold TH1 is associated with a first bin B(TH1). The time period between the first time period threshold TH1 and the second time period threshold TH2 is associated with a second bin B(TH2). The time period greater than the second time period threshold TH2 is associated with a third bin B(+). The occurrences of the plurality of demand time intervals DMT(N) are recorded in the respective bins as counts. Counts from the first and third bins are combined into a first count C(TH1) associated with setting the delay time interval 16 to be equal to about the first time period threshold. Counts from the second bin are recorded in a second count C(TH2) associated with setting the delay time interval to be equal to about the second time period threshold.

The delay time interval 16 is set based on a count having a majority, e.g., 4 counts of 7

1 samples, N. As an example, if the delay time interval is currently set to be equal to about the
2 first time period threshold TH1, and the measured demand time interval occurrences are recorded
3 in the second count C(TH2), then a majority of the delay time interval is set to be equal to about
4 the second time period threshold TH1. The delay time interval will be set back to be equal to
5 about the second time threshold when a majority of demand time interval occurrences are
6 recorded in the first count C(TH1).

7 Accordingly, recent past disk access patterns are used to predict future access patterns
8 and to manage and optimize the disk drive's power usage. The improved predictive power
9 management and usage techniques of the invention may extend the battery's power delivery time
10 thus providing longer operation of the mobile device between battery recharge operations.

11 With reference to FIG. 4, the disk drive 10 includes a disk control system 40 and a head
12 disk assembly (HDA) 42. The HDA includes a magnetic disk 44 having a disk media with
13 plurality of concentric tracks recorded thereon, a head 46 for writing user data to or reading user
14 data from a selected one of the tracks in response to a host command during operation of the disk
15 drive, and an actuator 48 for positioning the head over the selected track. The head in present
16 disk drives is a giant magneto-resistive (GMR) read element and a thin film inductive write
17 element. The actuator is typically implemented as a voice coil motor (VCM) which rotates the
18 actuator arm 50 about a pivot 52 in order to position the head radially over the disk in response to
19 control signals from the disk control system. The HDA includes a preamplifier 54, a spindle
20 motor 56 for rotating the disk, and a ramp 58 for parking the actuator arm and head when the
21 disk drive is in a spin-down mode. The head communicates with the disk control system via the
22 preamplifier. The preamplifier provides an amplified signal to a read/write channel 60 of disk
23 control system.

24 The disk control system 40 implements a servo control loop which causes the head 46 to
25 follow a centerline of a selected track in an operation generally referred to as "track following."
26 A servo processor 62 commands a servo controller 64 to control the position of the head over a
27 target track for subsequent execution of read or write commands. The servo processor receives a
28 representative form of a position signal sensed by the head, via the preamplifier 54 and the
29 read/write channel 60, and performs calculations to calculate a position error signal. The servo
30 controller responds to digital commands from the servo processor responsive to the position error

1 signal to provide a corresponding analog signal to the VCM. The servo processor also provides
2 commands to the servo controller to maintain a substantially constant operational rotational
3 velocity of the spindle motor 56.

4 The disk control system 40 also includes circuitry and processors that provide an
5 intelligent disk control system interface between a host processor 68 of the mobile device 12 and
6 the HDA 42 for execution of read and write commands. The disk control system includes an
7 integrated circuit host interface and disk controller (HIDC) 74 having a host interface 76, a buffer
8 controller 78 and a disk controller 80. The host interface 76 communicates with the host
9 processor. The buffer controller controls a buffer memory 82 employed for storing data from the
10 host processor that is to be written to the disk 44 and for storing data read from the disk. The
11 disk controller sends data to and receives data from the read/write channel 60 and provides for
12 error correction and error detection on data read from the disk. An interface processor 84
13 handles the flow of data commands received by the host interface by sending commands to and
14 reading status from the disk controller. The interface processor and the servo processor 62 may
15 be advantageously implemented using a single processor of sufficient processing power.
16 Information may be stored in nonvolatile memory 86 or in random access memory (RAM) 88
17 such as dynamic RAM (DRAM).

18 The method steps (FIG. 1) of the invention may be implemented by the interface
19 processor 84 and the HIDC 74. Program code for the steps may be stored in the nonvolatile
20 memory 86 and transferred to the RAM 88 during operation for execution by the interface
21 processor.

22 With reference to FIG. 5, the disk drive 10 has several power states or modes based on
23 usage patterns. The power states are active, idle1, idle2, idle3, and standby. In the active state,
24 all of the disk drive's electronic components are at full power, the disk 44 is spinning at its
25 standard rate, the actuator arm 50 is over the disk media of the disk, and a seek, read or write
26 operation is in progress. In the idle1 state, the disk is spinning at its standard rate, the actuator
27 arm is over the disk media, and the read/write channel 60 is turned off between servo samples.
28 In the idle2 state, the spindle motor 56 is turned off, the disk is not spinning, the actuator arm is
29 parked on the ramp 58 and latched, and the read/write channel and the servo controller 64 are in a
30 low power mode. In the idle3 state, the disk is spinning at a fraction of its standard active state

rate, the actuator arm is parked and latched, and the read/write channel and the servo controller are in a low power mode. The idle3 state may be an alternative or precursor to the idle2 state. In the standby state, the spindle motor is turned off and the disk is not spinning, the actuator is parked and latched, the read/write channel and the servo controller are in a low power mode, and clock signals to the microprocessors are inhibited. The DRAM is still active in the standby state.

Continuous operation of the spindle motor 56 in the idle 1 state consumes battery power from the mobile device. Advantageously, during a period of low disk activity, the disk drive 10 is transitioned to the idle2 state. The transition to the idle2 state includes removing a motor drive current from the spindle motor 56 in an operation generally referred to as “spinning down” the disk 44. Returning the disk drive to the active or the idle1 states involves returning the disk to its standard active state rate in an operation generally referred to as “spinning up” the disk.

Spinning up the disk consumes battery power at an increased rate until a steady state rotational rate is obtained. This transitional quantity of power 92 is an inherent power penalty associated with spinning down and then spinning up the disk. Accordingly, unnecessarily spinning down the disk may consume more battery power than maintaining the disk drive in the idle1 state. Efficient battery power management requires a balance between keeping the disk drive in a lower power mode versus transitioning it back to a higher power mode to service a disk access.

With reference to FIGS. 6 and 7, an exchange rate interval EX is defined as a time interval during which the energy that may be consumed by disk drive 10 in the idle1 power state is about equal to the energy consumed to transition the disk drive from the idle2 state to the idle1 state. An exemplary exchange rate interval for a disk drive for use in a mobile device may be about 3 seconds. The present invention may be embodied in a method to optimize power usage using three time thresholds. In this example, the first threshold TH1 may be 10 milliseconds, the second threshold TH2 may be 1 second, and the third threshold TH3 may be 2 seconds. The demand time intervals may be tracked by binning the measured demand time intervals into time periods between the thresholds. A first bin B(TH1) may be associated with a time period falling between 0 and 10 milliseconds. A second bin B(TH2) may be associated with a time period falling between 10 milliseconds and 1 second. A third bin B(TH3) may be associated with a time period falling between 1 second and 2 seconds. A fourth bin B(G), for a grey interval, may be associated with a time period between the third threshold TH3, e.g. 2 seconds, and the

exchange rate interval EX, e.g., 3 seconds. The grey interval is close to the exchange rate interval and the potential power savings benefits are relatively small if disk access commands are arriving in this interval. However, the risk of wasting energy by staying in the idle1 state to service disk access commands arriving in the grey interval is considerable. For example, if the disk drive is transitioned from the idle 1 state to idle2 state about 3 seconds after servicing a host command, and a subsequent host command is received after 3.5 seconds, then the 3 second of operation in the idle 1 mode is wasted. Battery power may have been used more effectively if the disk drive was transitioned to the idle2 state after 10 milliseconds. However, if the disk drive is scheduled for transition from the idle 1 state to the idle2 state after about 3 seconds after servicing a host command, and a subsequent command is received after 2.7 seconds, then the power benefit is related to 0.3 second of operation. Accordingly, the cost (3 seconds) is much greater than the benefit (0.3 seconds). A fifth bin B(EX+), for a dark interval, may be associated with demand time period greater than the exchange rate interval EX.

In another exemplary embodiment assuming a 3 second exchange rate interval and employing only two thresholds, the first threshold may be 10 milliseconds and the second threshold may be 2 seconds.

Counts from the first bin B(TH1), the fourth bin B(G), and the fifth bin B(EX+) may be combined into a first count C(TH1) associated with setting the delay time interval 16 to be equal to about the first time period threshold TH1. Counts from the second bin B(TH2) may be recorded in a second count C(TH2) associated with setting the delay time interval to be equal to about the second time period threshold TH2. Counts from the third bin B(TH3) may be recorded in a third count C(TH3) associated with setting the delay time interval to be equal to about the third time period threshold TH3.

The delay time interval 16 is set based on the count having a majority. In FIG. 7, the first count C(TH1) has a majority (4 counts out of 7 samples N). Accordingly, the delay time interval is set to be equal to about the first time interval threshold TH1. If no count has a majority, then the delay time interval remains at its current setting. For example, for a sample size of 7 and a majority limit of 4, if the first count C(TH1) is 3, the second count C(TH2) is 2, and the third count C(TH3) is 2, then none of the counts reaches the majority limit and the delay time interval remains unchanged.

- 1
- 2
- 3